Video denoising using 2D and 3D dual-tree complex wavelet transforms

Ivan W. Selesnick and Ke Yong Li

Polytechnic University, 6 Metrotech Center, Brooklyn, New York 11201

ABSTRACT

The denoising of video data should take into account both temporal and spatial dimensions, however, true 3D transforms are rarely used for video denoising. Separable 3-D transforms have artifacts that degrade their performance in applications. This paper describes the design and application of the non-separable oriented 3-D dual-tree wavelet transform for video denoising. This transform gives a motion-based multi-scale decomposition for video — it isolates in its subbands motion along different directions. In addition, we investigate the denoising of video using the 2-D and 3-D dual-tree oriented wavelet transforms, where the 2-D transform is applied to each frame individually.

Keywords: Complex wavelet transform, denoising, oriented, dual-tree, 3D wavelet

1. INTRODUCTION

Although video is a 3-D data set, 3-D data transforms are usually not used for its representation — indeed, the standard 3-D data transforms do not provide useful representations for most video data. (By useful, we mean that the transform should have a good energy compaction property and that the transform coefficients reflect properties of the data in which we are interested.) Note that most 3-D data transforms are separable implementations of 1-D transforms. However, separable implementations of multidimensional (M-D) transforms have artifacts that can seriously reduce their usefulness, especially for the representation of M-D image data (natural images, video, medical volume data, etc). For example, the standard separable 3-D transform is rarely used for video representation because it mixes 3-D orientations in its subbands. This artifact, which is illustrated below and in accompanying video clips on the web, prevents the transform from providing an effective representation of video. In order for a 3-D data transform to provide a useful representation for M-D image data, the transform probably has to be non-separable. However, non-separable implementations are more computationally expensive; and in addition, they are often more difficult to design.

It turns out, that a recently developed type of complex wavelet transform can be used to implement fast non-separable real-valued multidimensional wavelet transforms that overcome the serious artifacts of the separable M-D wavelet transform. In this paper, we propose to use the 3-D *dual-tree* wavelet transform for video noise reduction.

Although the dual-tree transform is expansive (it expands an N-sample data vector to M wavelet coefficients with $M > N$), it is free of the mixing artifact and gives a meaningful multi-scale decomposition for video. With the dual-tree transform, it is more likely that the multiresolution framework, which has proven very effective for image denoising can also be effectively applied to video denoising. The dual-tree transform isolates motion in different directions in separate subbands; the direction of motion can be inferred to some degree from the wavelet coefficients. Because this oriented 3-D transform can represent motion information, it provides a tool for video denoising that takes into account the motion of image elements, without explicitly using motion estimation.

Programs for 3-D oriented wavelet transforms and examples will be available on the web at `http://taco.poly.edu/WaveletSoftware/`.

Further author information: (Send correspondence to I.W.S.)
I.W.S: E-mail: selesi@poly.edu, Telephone: (718) 260-3416.
K.Y.L.: E-mail: kli10@utopia.poly.edu

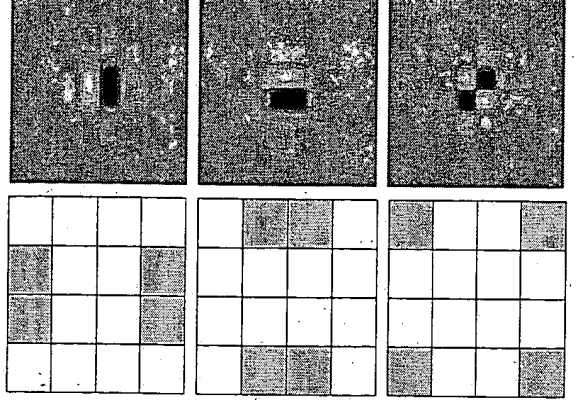

Figure 1. Typical wavelets associated with the 2-D separable discrete wavelet transform. The top row illustrates the wavelets in the spatial domain, the second row illustrates the (idealized) support of the spectrum of each wavelet in the 2-D frequency plane. The checkerboard artifact of the third wavelets is apparent in the frequency domain as well as the spatial domain.

2. MULTIDIMENSIONAL WAVELET TRANSFORMS

In a sense, wavelet-bases are optimal for a large class of one dimensional signals (including many real signals). However, as has been recognized, the (separable) 2-D wavelet transform does not possess these optimality properties for natural images.[1] The reason for this is that while the separable 2-D wavelet transform represents point-singularities efficiently, it is *inefficient* for line- and curve-singularities (edges). (Even so, it is better than the DCT for this — which explains the JPEG-2000 compression standard.)

Some of the important developments in recent wavelet research has been the implementation of 2-D multiscale transforms that represent edges more efficiently than does the separable wavelet transform. Examples include curvelets,[2,3] directional filter banks and pyramids,[4,5] complex filter banks,[6] the steerable pyramid,[7,8] and the complex dual-tree wavelet transform.[9,10] These transforms give superior results for image processing applications compared to the separable wavelet transform. In this paper, we investigate the use of the 3-D version of the dual-tree complex wavelet transform for video noise reduction. The dual-tree wavelet transform is nearly-shift invariant, isolates edges with different orientations in different subbands, and has a manageable redundancy. It will also be of interest to extend to 3-D, the curvelet transform, steerable pyramid, and other oriented multiscale transforms[3–6]; and to investigate their use for video denoising.

To understand why the dual-tree wavelet transform outperforms the separable wavelet transform, it is useful to compare the wavelets associated with the two transforms. They are illustrated in Figures 1 and 2 where it is shown that the separable 2-D wavelet transform is characterized by three wavelets, the dual-tree 2-D transform by six wavelets. The third wavelet associated with the separable wavelet transform has a *checkerboard* appearance — because it mixes the 45 and -45 degree orientations; equivalently, *the separable transform fails to isolate these orientations.* On the other hand, the dual-tree transform succeeds in isolating different orientations — each of the six wavelets are aligned along a specific direction and no checkerboard effect appears. Moreover, the dual-tree transform covers more distinct orientations than does the separable transform.

One way to understand the checkerboard artifact that arises in the separable implementation of the 2-D wavelet transform is in the frequency domain. If $\psi(x)$ is a real wavelet and the 2-D separable wavelet is given by $\psi(x,y) = \psi(x)\,\psi(y)$, then the spectrum of $\psi(x,y)$ is illustrated by the following idealized diagram.

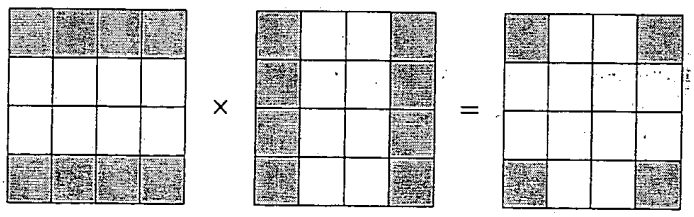

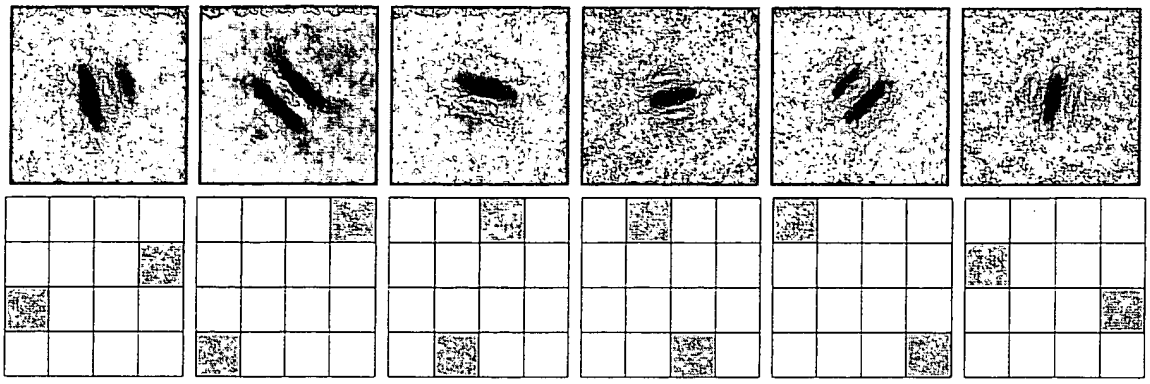

Figure 2. Typical wavelets associated with the 2-D dual-tree oriented wavelet transform. The top row illustrates the wavelets in the spatial domain, the second row illustrates the (idealized) support of the spectrum of each wavelet in the 2-D frequency plane. The absence of the checkerboard phenomenon is observed in both the spatial and frequency domains.

Because $\psi(x)$ is a real function its spectrum must be two-sided; therefore, the checkerboard artifact is unavoidable in the frequency domain. Likewise, the checkerboard artifact arises in the spatial domain as well.

To explain how the dual-tree complex wavelet transform produces oriented wavelets, consider the 2-D wavelet $\psi(x,y) = \psi(x)\,\psi(y)$ associated with the row-column implementation of the DWT, where $\psi(x)$ is a complex (approximately analytic) wavelet given by $\psi(x) = \psi_h(x) + j\,\psi_g(x)$. We obtain for $\psi(x,y)$ the expression,

$$\psi(x,y) = [\psi_h(x) + j\,\psi_g(x)]\,[\psi_h(y) + j\,\psi_g(y)] \tag{1}$$

$$= \psi_h(x)\,\psi_h(y) - \psi_g(x)\,\psi_g(y) + j\,[\psi_g(x)\,\psi_h(y) + \psi_h(x)\,\psi_g(y)] \tag{2}$$

The support of the spectrum of this complex wavelet is illustrated by the following idealized diagram.

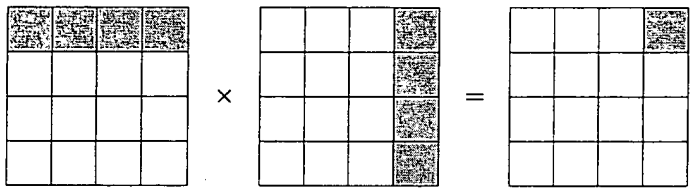

Because the spectrum of the (approximately) analytic 1-D wavelet is supported on only one side of the frequency axis, the spectrum of the complex 2-D wavelet $\psi(x,y)$ is supported in only one quadrant of the 2-D frequency plane. For this reason, the complex 2-D wavelet is oriented in a specific direction. If we take the real part of this complex wavelet, we obtain the sum of two separable wavelets:

$$\text{Real Part}\{\psi(x,y)\} = \psi_h(x)\,\psi_h(y) - \psi_g(x)\,\psi_g(y) \tag{3}$$

Because the spectrum of a real function must be symmetric with respect to the origin, the spectrum of this real wavelet is supported in the two quadrants of the 2-D frequency plane, as illustrated in the following (idealized) diagram.

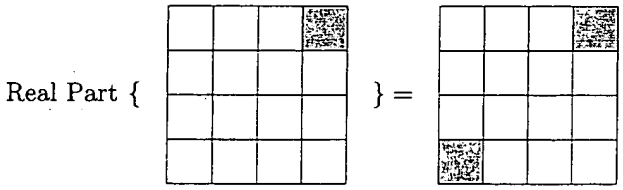

Unlike the real separable wavelet, the support of the spectrum of this real wavelet does not posses the checkerboard artifact and therefore the real wavelet is oriented. Note that this construction depends on the complex wavelet $\psi(x) = \psi_h(x) + j\,\psi_g(x)$ being (approximately) analytic; or equivalently, that $\psi_g(t)$ is approximately the Hilbert transform of $\psi_h(t)$, [$\psi_g(t) \approx \mathcal{H}\{\psi_h(t)\}$].

If $\psi_g(t) \approx \mathcal{H}\{\psi_h(t)\}$, then the same derivation can be used to implement an oriented nonseparable 2-D wavelet transform by combining the subbands of two separable 2-D DWTs. If we define two separable 2-D wavelet bases in the usual manner,

$$\psi_{1,1}(x,y) = \phi_h(x)\,\psi_h(y) \qquad \psi_{2,1}(x,y) = \phi_g(x)\,\psi_g(y) \tag{4}$$

$$\psi_{1,2}(x,y) = \psi_h(x)\,\phi_h(y) \qquad \psi_{2,2}(x,y) = \psi_g(x)\,\phi_g(y) \tag{5}$$

$$\psi_{1,3}(x,y) = \psi_h(x)\,\psi_h(y) \qquad \psi_{2,3}(x,y) = \psi_g(x)\,\psi_g(y). \tag{6}$$

then the 6 wavelets defined by

$$\psi_i(x,y) = \frac{1}{\sqrt{2}}\left(\psi_{1,i}(x,y) + \psi_{2,i}(x,y)\right) \tag{7}$$

$$\psi_{i+3}(x,y) = \frac{1}{\sqrt{2}}\left(\psi_{1,i}(x,y) - \psi_{2,i}(x,y)\right) \tag{8}$$

for $1 \le i \le 3$ are oriented, as illustrated in Figure 2. Because the sum/difference operation is orthonormal, this constitutes a perfect reconstruction wavelet transform. (If the 1-D wavelets $\psi_g(t)$ and $\psi_h(t)$ form orthonormal bases, then this constitutes a *tight frame*, or a *self-inverting* transform.)

In 3-D, the checkerboard artifact of the separable transform is more serious than in 2-D. Correspondingly, the gain provided by using the oriented wavelet transform in place of the separable one is greater in higher dimensions.

2.1. Developing the Fast Oriented 2-D Wavelet Transform

The dual-tree wavelet transform is implemented using separable transforms and by combining subband signals appropriately. So even though it is non-separable (and therefore free of some of the limitations of separable transforms) it inherits the computational efficiency of separable transforms. Specifically, the 1-D dual-tree wavelet transform is implemented using two filter banks in parallel operating on the same data as illustrated in Figure 3. A complex-valued wavelet $\psi(t)$ can be obtained as $\psi(t) = \psi_h(t) + j\,\psi_g(t)$ where $\psi_h(t)$ and $\psi_g(t)$ are both real-valued wavelets. To implement directional 2-D transforms using the dual-tree transform, it is necessary that the spectrum of the complex-valued wavelet $\psi(t)$ be single sided ($\Psi(\omega) = 0$ for $\omega \le 0$). It is therefore required that the wavelet $\psi_g(t)$ be the Hilbert transform of the wavelet $\psi_h(t)$,

$$\Psi_g(\omega) = \begin{cases} -j\,\Psi_h(\omega), & \omega > 0 \\ j\,\Psi_h(\omega), & \omega < 0 \end{cases} \tag{9}$$

which we denote as $\psi_g(t) = \mathcal{H}\{\psi_h(t)\}$. But how does one design and implement two discrete wavelet transforms so that the wavelets associated with them form a Hilbert pair? Specifically, how should the two lowpass CQF filters $h_0(n)$ and $g_0(n)$ in Figure 3 satisfying perfect reconstruction properties be designed so that $\psi_g(t) = \mathcal{H}\{\psi_h(t)\}$? These questions have been addressed in our recent work.[11, 12]

2.2. Hilbert Pairs of Wavelet Bases: Characterization Theorem

Hilbert pairs of wavelet bases were studied in[11, 12] which give a characterization and provides a Daubechies-like construction for approximate Hilbert pairs of orthonormal (and biorthogonal) wavelets with vanishing moments and compact support. To outline these results, let the real filters $h_0(n)$, $h_1(n)$ represent a conjugate quadrature filter (CQF) pair.[13] That is, the autocorrelation $p_h(n)$ of the lowpass filter $h_0(n)$ is halfband: $p_h(2\,n) = \delta(n)$ where $p_h(n) := h_0(n) * h_0(-n)$; and the highpass filter is given by $h_1(n) = (-1)^n\,h_0(M-n)$ where M is an odd integer. Let the real filters $g_0(n)$, $g_1(n)$ represent a second CQF pair. The real-valued scaling function $\phi_h(t)$ and

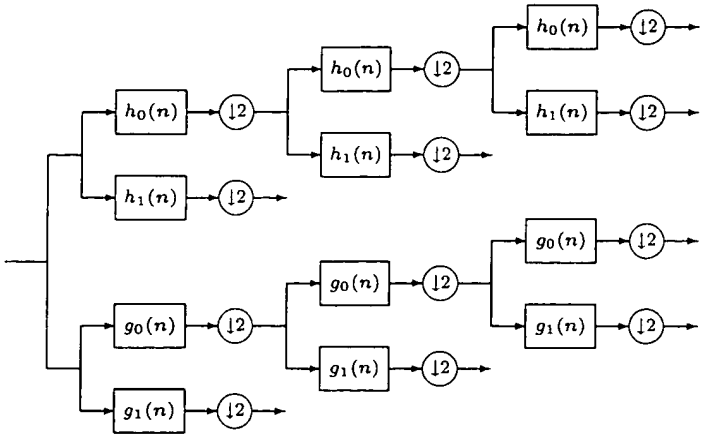

Figure 3. The 1-D dual-tree wavelet transform is implemented using a pair of filter banks operating on the same data simultaneously. The upper iterated filter bank represents the real part of a complex wavelet transform. The lower one represents the imaginary part. The transform is an expansive (or oversampled) transform (or *frame*).

the real-valued wavelet $\psi_h(t)$ associated with the pair $(h_0(n), h_1(n))$ are defined implicitly by the following pair of equations:

$$\phi_h(t) = \sqrt{2} \sum_n h_0(n)\, \phi_h(2t-n), \qquad \psi_h(t) = \sqrt{2} \sum_n h_1(n)\, \phi_h(2t-n). \tag{10}$$

The real-valued scaling function $\phi_g(t)$ and wavelet $\psi_g(t)$ associated with the pair $(g_0(n), g_1(n))$ are defined similarly. Note that if $|H_0(e^{j\,\omega})| = |G_0(e^{j\,\omega})|$ then it can be shown that $|\Psi_h(\omega)| = |\Psi_g(\omega)|$. Therefore, the two lowpass filters should be related as

$$G_0(\omega) = H_0(\omega)\, e^{-j\,\theta(\omega)} \tag{11}$$

where $\theta(\omega)$ is 2π-periodic. But how should the phase function $\theta(\omega)$ be chosen so that the two wavelets generated by $h_0(n)$ and $g_0(n)$ form a Hilbert transform pair? Using the infinite-product formula it can be shown that (11) leads to the following relationship between the wavelet spectrums:

$$\Psi_g(\omega) = \Psi_h(\omega)\, \exp\left(j \left[\theta(\omega/2 - \pi) - \sum_{k=2}^{\infty} \theta(\omega/2^k) \right] \right). \tag{12}$$

From (9) and (12) we see that $\theta(\omega)$ must satisfy the condition:

$$\theta(\omega/2 - \pi) - \sum_{k=2}^{\infty} \theta(\omega/2^k) = \begin{cases} -\frac{\pi}{2}, & \omega > 0 \\ \frac{\pi}{2}, & \omega < 0. \end{cases} \tag{13}$$

We have shown in[11] that if $\theta(\omega)$ is given by $\theta(\omega) = 0.5\,\omega$ for $|\omega| < \pi$ then condition (13) holds and $\psi_g(t) = \mathcal{H}\{\psi_h(t)\}$.

Theorem: *If $H_0(e^{j\,\omega})$ and $G_0(e^{j\,\omega})$ are lowpass CQF filters with*

$$G_0(e^{j\,\omega}) = H_0(e^{j\,\omega})\, e^{-j\,0.5\,\omega} \quad \textit{for} \quad |\omega| < \pi, \tag{14}$$

then the corresponding wavelets are a Hilbert transform pair, $\psi_g(t) = \mathcal{H}\{\psi_h(t)\}$.

Equivalently, the digital filter $g_0(n)$ is a *half-sample* delayed version of $h_0(n)$, $g_0(n) = h_0(n-0.5)$. If it is desired that both lowpass filters $h_0(n)$ and $g_0(n)$ be FIR, then it is impossible to satisfy (14) exactly and an approximation must be made. Informally, if in the time-domain we have $g_0(n) \approx h_0(n-0.5)$, then $\psi_g(t) \approx \mathcal{H}\{\psi_h(t)\}$. In[12] we have developed a Daubechies-like algorithm for the construction of Hilbert pairs of short orthonormal (and

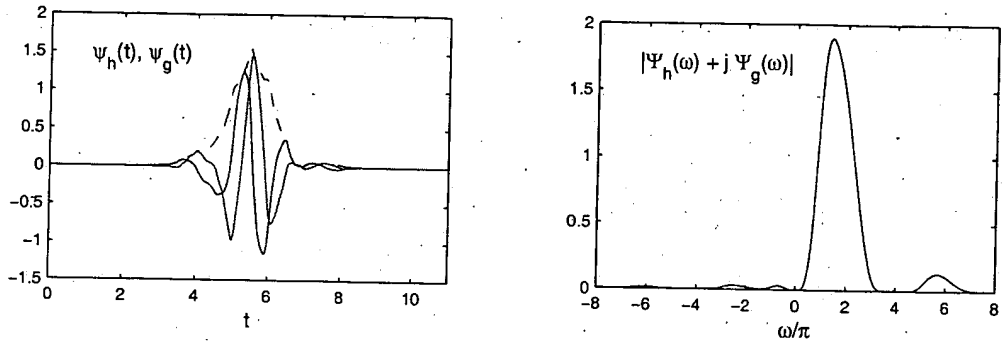

Figure 4. Approximate Hilbert transform pair of orthonormal wavelet bases designed according to.[12] Left: The wavelets $\psi_h(t)$, $\psi_g(t)$ are shown by blue and green solid lines. The magnitude of the complex wavelet $\psi(t) = \psi_h(t) + j\,\psi_g(t)$ is shown as a dashed line. Right: The spectrum of the complex wavelet $\psi(t)$ is single sided. The directional 2-D wavelet transform (Figure 2) and the motion-selective 3-D wavelet transform are efficiently implemented using Hilbert pairs of 1-D wavelets like those illustrated here.

biorthogonal) wavelet bases. As an example, a Hilbert pair of wavelets is illustrated in Figure 4. This construction yields pairs of bases which can be used to efficiently implement directional 2-D wavelet transforms, the associated wavelets of which are illustrated in Figure 2. In addition, these pairs of filters can be used to efficiently implement the motion-selective wavelet transform.

3. MOTION-BASED WAVELET TRANSFORM FOR VIDEO REPRESENTATION

The shortcoming of the separable 2-D wavelet transform for image processing is further compounded for video processing — with multidimensional separable transforms even more mixing of different orientations occurs. (Note that with video, an object edge moving in a particular spatial orientation contributes to a 3-D edge.) For this reason, 3-D wavelet transforms have not been widely applied in video restoration. Previous works on video restoration[14–19] usually perform noise reduction along motion trajectories. In some algorithms, a separable wavelet transform is applied along the motion trajectory in the temporal direction: i.e. first determine the motion trajectory of a pixel or block, then process along that trajectory in the temporal direction. The performance, however, is dependent on the accuracy of motion estimation which is more difficult for noisy video.

The principle by which the 2-D dual-tree DWT resolves the problem of the mixing of orientations, can also be used to resolve the mixing of orientations in the 3-D case. This dual-tree wavelet transform is *motion-selective*, while the separable 3-D wavelet transform is not. (Likewise, the 2-D dual-tree transform is *direction-selective*, while the 2-D transform is not.)

Figures 5 and 6 illustrates the difference between the separable 3-D wavelet transform and the dual-tree 3-D transform. The figures depict a typical 3-D wavelet $\psi(x, y, t)$ associated with each of the two wavelet transforms. The isosurface provides a way to represent a function of three variables, similar to a contour plot of a function of two variables. To interpret the isosurfaces in Figure 5 note that, like a contour plot, the points on the curves (surfaces) are points where the functions are equal-valued. The wavelet associated with *the separable 3-D transform has the same checkerboard phenomenon present in the separable 2-D case*, a consequence of the mixing of orientations. The wavelet associated with *the dual-tree 3-D transform is free of this effect*, as illustrated in Figure 6 and the accompanying video clips. In addition, the dual-tree 3-D transform has many more subbands than the separable 3-D transform (28 subbands instead of 7). These 28 subbands capture motion having different directions and different velocities.

The video clips that accompany this paper (available on the web) reveal the behavior of these 3-D functions better than Figure 6 can. When the wavelets of the dual-tree 3-D transform are displayed as a movie, one observes a finite-time motion in a localized patch in a particular direction — the dual-tree 3-D wavelet transform therefore reveals motion and edge information. The subbands of the dual-tree 3-D wavelet transform can isolate 2-D edges with different orientations that are moving in different directions. The video clips show, on the other

hand, that the separable wavelet conveys/captures limited motion information and conveys/captures limited edge information in a video frame.

The idea of using a complex (analytic) wavelet transform to implement a real 3-D oriented wavelet transform for motion analysis was described originally by Burns *et al.*[20,21] The implementation described in these references also uses the Hilbert transform to avoid the checkerboard artifact. In these papers, the Hilbert transform is first applied to the data, the separable wavelet transform is applied to both the original data and the Hilbert transformed data, and then the wavelet coefficients are combined as in the dual-tree transform, to obtain an oriented transform. On the other hand, in the dual-tree implementation, the Hilbert transform is built into the wavelet transform itself.

Note that the ideal Hilbert transform is represented by an infinitely long impulse response that has slow decay. The use of the ideal (or near ideal) Hilbert transform in conjunction with the wavelet transform effectively increases the support of the wavelets. In order that the wavelets have short support, an approximate Hilbert transform that is more localized in time should be used instead. However, the accuracy of the approximate Hilbert transform should depend on the scale of the wavelet transform. When the Hilbert transform is applied first to the data, a single Hilbert transform is effecting wavelet coefficients at all scales, while when the Hilbert transform is built into wavelet transform as in the dual-tree implementation, the Hilbert transform scales with the wavelet coefficients.

The 3-D oriented dual-tree wavelet transform is developed similarly to the 2-D version. Suppose that $\psi_g(t)$ is approximately the Hilbert transform of $\psi_h(t)$, $[\psi_g(t) \approx \mathcal{H}\{\psi_h(t)\}]$. Define the approximately analytic complex wavelet $\psi(\cdot)$ by $\psi(x) = \psi_h(x) + j\,\psi_g(x)$. Then consider the complex 3-D wavelet,

$$\psi(x,y,z) = \psi(x)\,\psi(y)\,\psi(z) \tag{15}$$

$$= (\psi_h(x) + j\,\psi_g(x))\,(\psi_h(y) + j\,\psi_g(y))\,(\psi_h(z) + j\,\psi_g(z)). \tag{16}$$

As in the 2-D case, take the real part of $\psi(x,y,z)$ to get an oriented real 3-D wavelet:

$$\text{Real Part}\{\psi(x,y,z)\} = \psi_h(x)\,\psi_h(y)\,\psi_h(z) - \psi_g(x)\,\psi_g(y)\,\psi_h(z) - \psi_g(x)\,\psi_h(y)\,\psi_g(z) - \psi_h(x)\,\psi_g(y)\,\psi_g(z).$$

We can write it more compactly as

$$\text{Real Part}\{\psi(x,y,z)\} = \psi_1(x,y,z) - \psi_2(x,y,z) - \psi_3(x,y,z) - \psi_4(x,y,z)$$

where we define

$$\psi_1(x,y,z) := \psi_h(x)\,\psi_h(y)\,\psi_h(z) \tag{17}$$

$$\psi_2(x,y,z) := \psi_g(x)\,\psi_g(y)\,\psi_h(z) \tag{18}$$

$$\psi_3(x,y,z) := \psi_g(x)\,\psi_h(y)\,\psi_g(z) \tag{19}$$

$$\psi_4(x,y,z) := \psi_h(x)\,\psi_g(y)\,\psi_g(z). \tag{20}$$

In the 3-D case, it is necessary to combine four separable wavelet transforms, instead of two as in the 2-D case. To obtain the remaining subbands, we take in addition to (15), the real part of $\psi(x)\,\psi(y)\,\overline{\psi(z)}$, $\psi(x)\,\overline{\psi(y)}\,\psi(z)$, and $\psi(x)\,\overline{\psi(y)}\,\overline{\psi(z)}$, where the overline represents complex conjugation. This gives the following orthonormal combination matrix of the four separable 3-D wavelet transforms:

$$\psi_a(x,y,z) = 0.5\;(\psi_1(x,y,z) - \psi_2(x,y,z) - \psi_3(x,y,z) - \psi_4(x,y,z))$$
$$\psi_b(x,y,z) = 0.5\;(\psi_1(x,y,z) - \psi_2(x,y,z) + \psi_3(x,y,z) + \psi_4(x,y,z))$$
$$\psi_c(x,y,z) = 0.5\;(\psi_1(x,y,z) + \psi_2(x,y,z) - \psi_3(x,y,z) + \psi_4(x,y,z))$$
$$\psi_d(x,y,z) = 0.5\;(\psi_1(x,y,z) + \psi_2(x,y,z) + \psi_3(x,y,z) - \psi_4(x,y,z))$$

By applying this combination matrix to each of the seven subbands, one obtains the 3-D oriented dual-tree wavelet transform illustrated in Figure 6. A program implementing this oriented transform is available on the web.

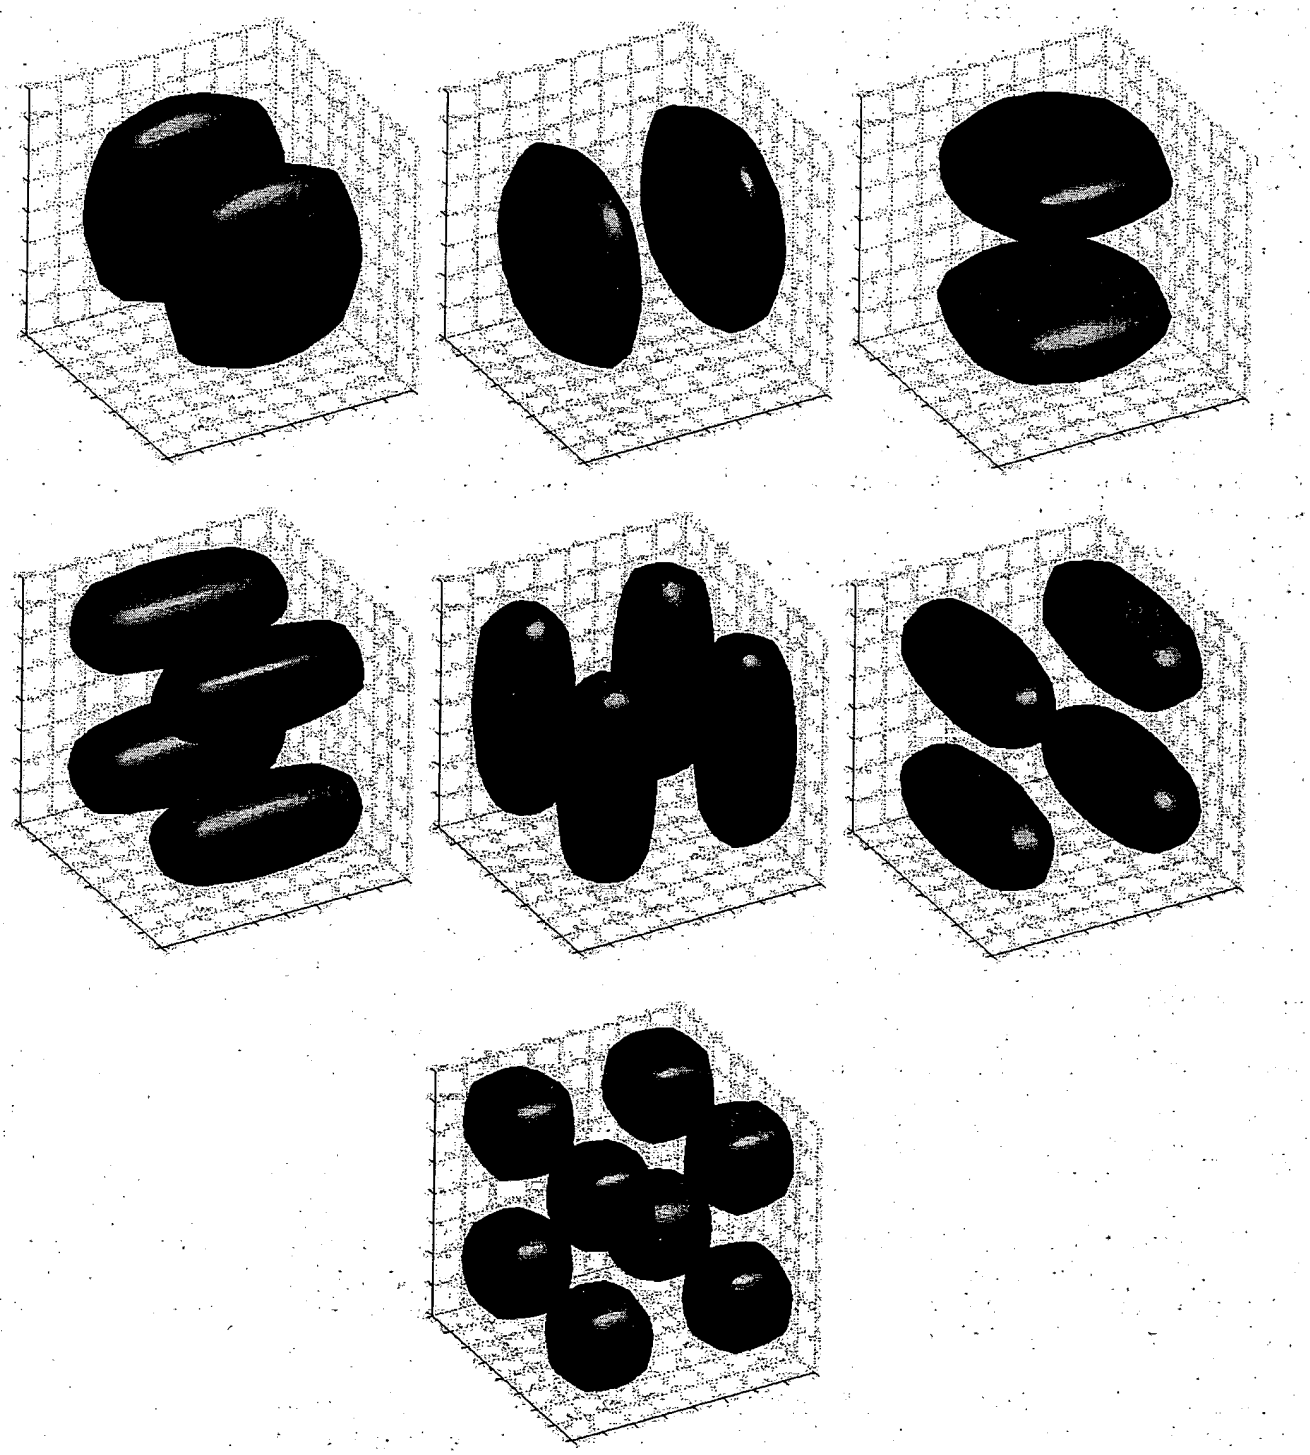

Figure 5. Isosurfaces of the 7 3-D wavelets for a typical separable 3-D wavelet transform. The blue and red surfaces represent the points where the wavelet $\psi(x, y, t)$ is equal-valued (blue: $\psi(x, y, t) = 0.3$; red: $\psi(x, y, t) = -0.3$).

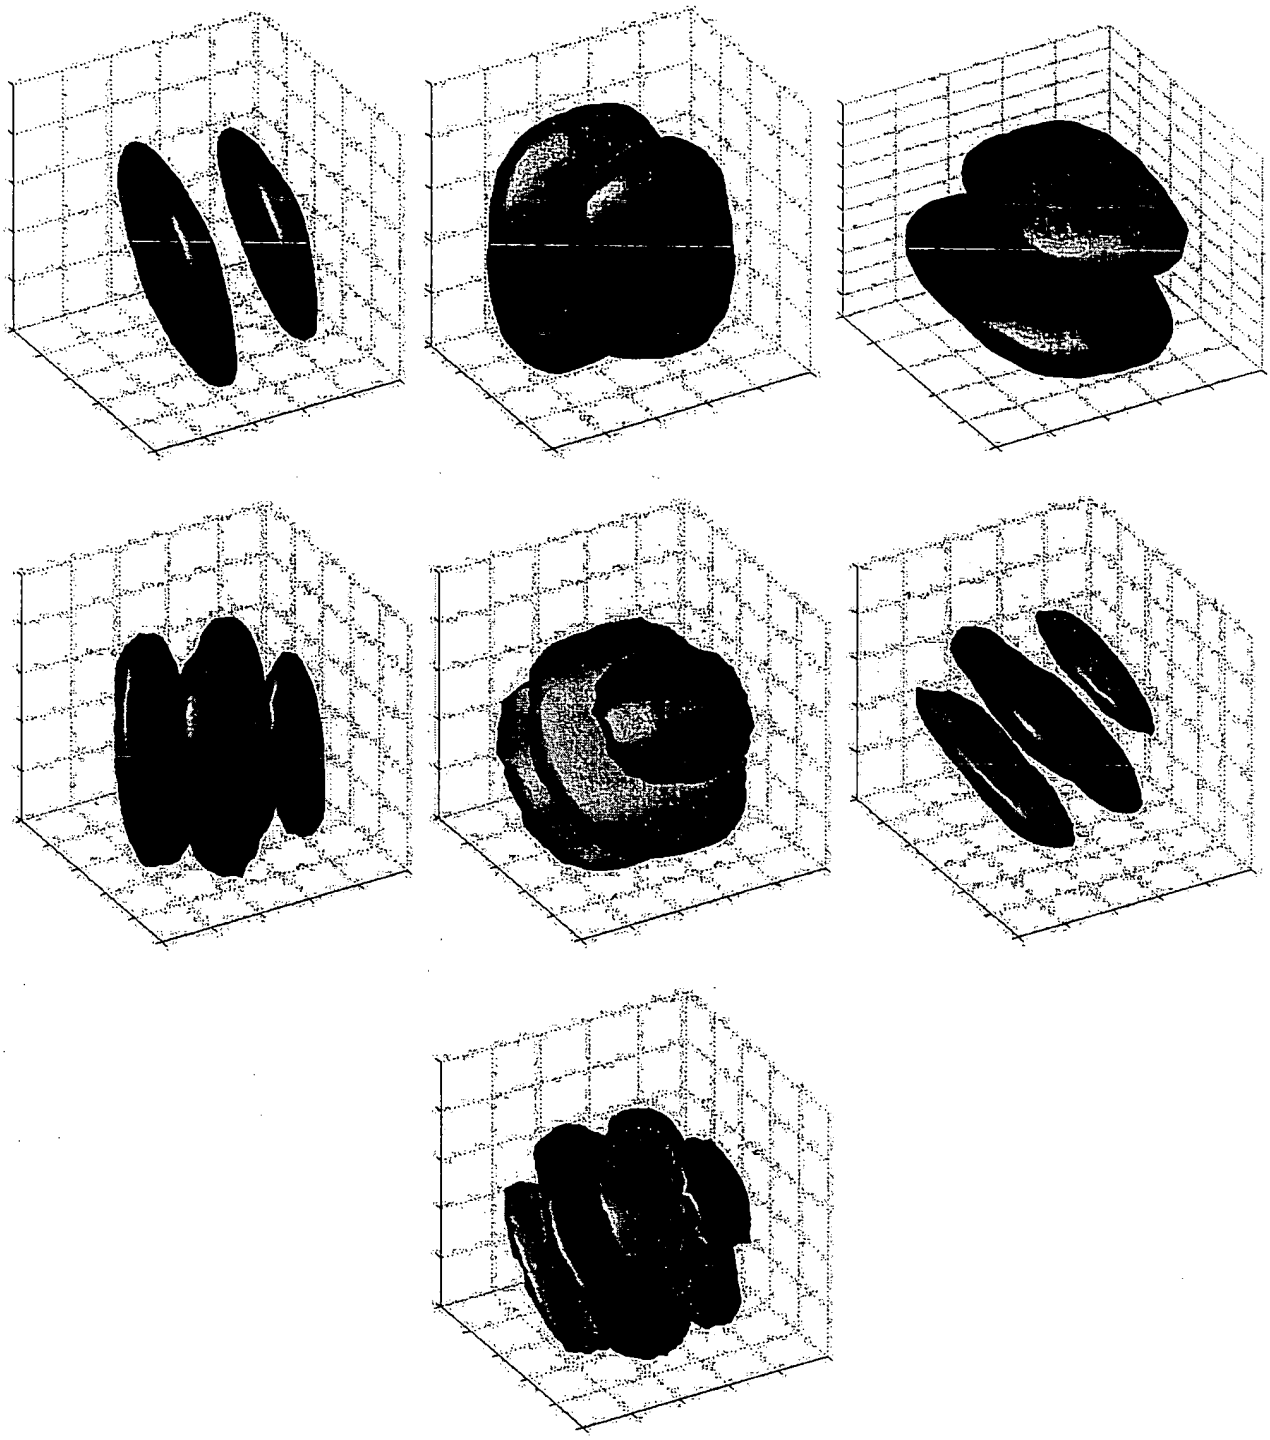

Figure 6. Isosurfaces of 7 of the 28 3-D wavelets for a typical oriented 3-D dual-tree wavelet transform. For the dual-tree 3-D wavelet transform, each subband corresponds to motion in a specific direction. Video clips available on the web display the wavelets as evolving gray-scale images.

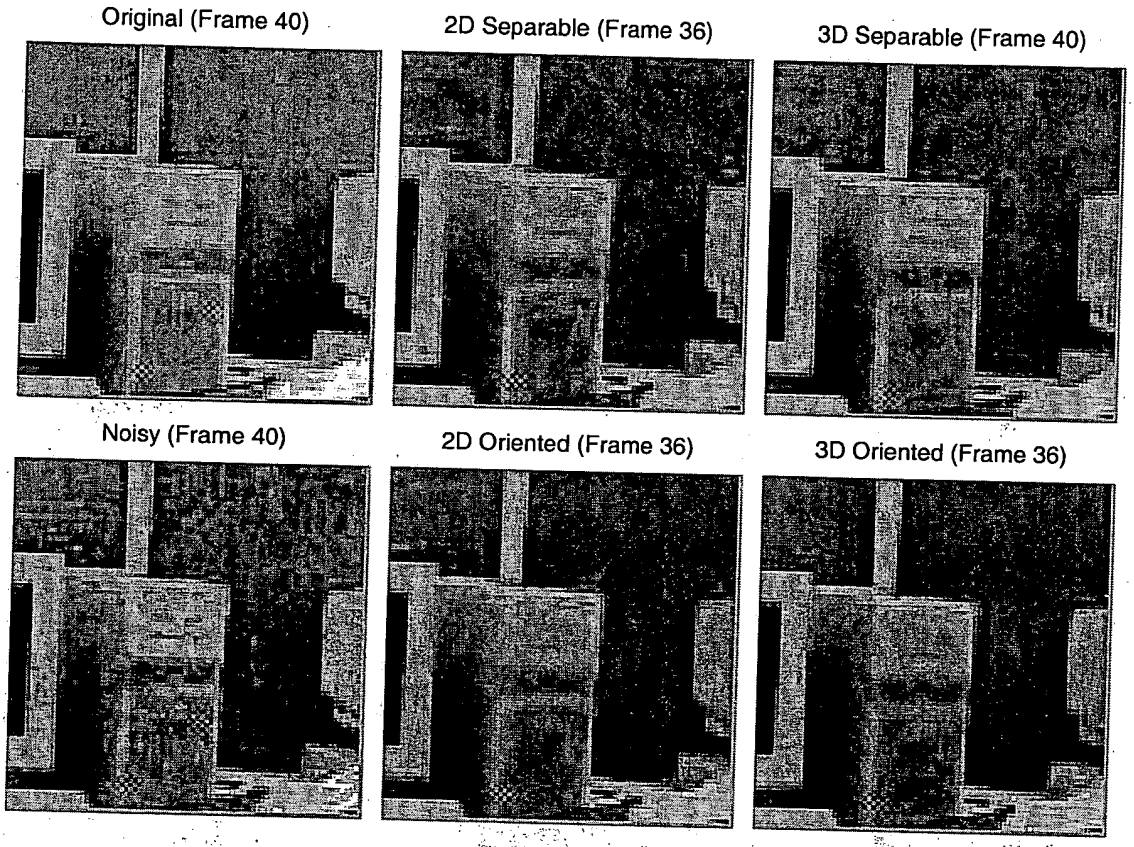

Figure 7. Denoising results for frame 36 of the test sequence.

4. DENOISING EXAMPLE

To evaluate the use of the 3-D oriented dual-tree wavelet transform for video denoising, the soft thresholding scheme was used with four different transforms. In the first two cases, we applied the 2-D separable and the 2-D oriented wavelet transforms to each frame of the video. In the second two cases, we applied the 3-D separable and the 3-D oriented wavelet transform to the 3-D video data. For the purpose of comparison, for each of the four transforms the optimal threshold was found using the original noise-free data. The same threshold was used for each subband of the respective types of wavelet transform.

The test video consists of a stationary view with a person walking across the scene. In the example, the root-mean-square (RMS) error for the separable 2-D transform was 7.6; for the dual-tree 2-D transform it was 6.5; for the separable 3-D transform it was 6.65; for the dual-tree 3-D transform it was 4.8. The dual-tree 3-D transform performed best for this example. Frames from the test video, the noisy video, and each of the processed videos are shown in Figure 7. The improvement of the 3-D dual-tree over the 2-D dual-tree is not clearly visible in the figure, however, the difference between them is more visible when viewing the video proper.

The oriented 3-D dual-tree transform does have its limitations however. For frames containing fast motion, the dual-tree 2-D transform can give a superior result. Figure 8 shows frames for each of the videos, containing fast motion: a person is passing in the scene. For this frame, the 2-D dual-tree transform does best because for fast motion it is more difficult to exploit the temporal correlation of pixel values. The 2-D and 3-D oriented transforms each have their benefits for video denoising. For the problem of image denoising, Starck, Candés and Donoho demonstrated[22] the improvement gained by using both the curvelet and the separable wavelet transforms together, with respect to either alone. In our current work we are investigating likewise the combined use of the 2-D and 3-D oriented wavelet transforms for video denoising.

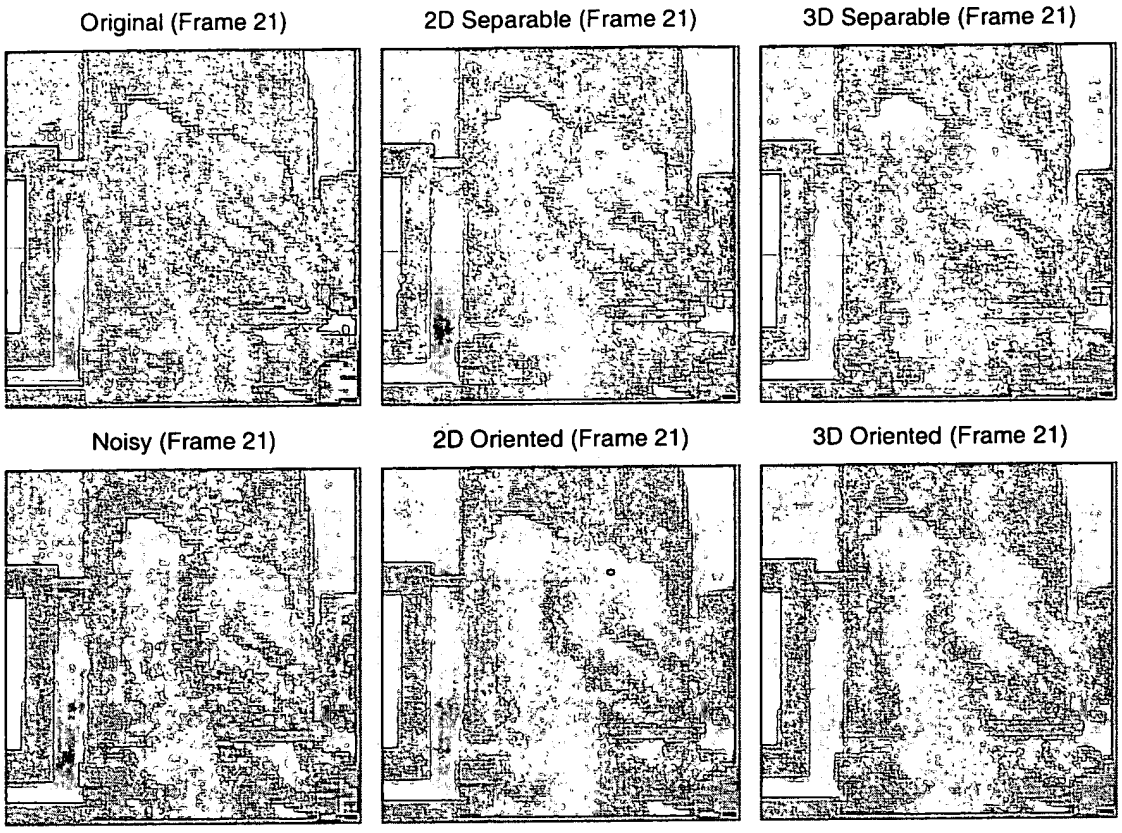

Figure 8. Denoising results for frame 21 of the test sequence.

ACKNOWLEDGMENTS

We gratefully acknowledge support from ONR (ONR grant N00014-03-1-0217), from NSF (NSF REU grant EEC-0139049, "Research Experience in Information Systems for Undergraduates") and from the Othmer Institute for Interdisciplinary Studies at Polytechnic University.

REFERENCES

1. D. L. Donoho, "Unconditional bases are optimal for data compression and for statistical estimation," *Applied and Computational Harmonic Analysis* **1**, pp. 100–115, Dec. 1993.
2. E. J. Candés and D. L. Donoho, "Curvelets, multiresolution representation, and scaling laws," in *Proceedings of SPIE*, **4119**, (San Diego), July 2000. Wavelet Applications in Signal and Image Processing VIII.
3. M. N. Do and M. Vetterli, "Pyramidal directional filter banks and curvelets," in *Proc. IEEE Int. Conf. Image Processing*, Oct. 2001.
4. R. H. Bamberger and M. J. T. Smith, "A filter bank for the directional decomposition of images: Theory and design," *IEEE Trans. on Signal Processing* **40**, pp. 882–892, Apr. 1992.
5. M. N. Do and M. Vetterli, "Contourlets: A directional multiresolution image representation," in *Proc. IEEE Int. Conf. Image Processing*, 2002.
6. F. Fernandes, R. van Spaendonck, M. Coates, and S. Burrus, "Directional complex-wavelet processing," in *Proceedings of SPIE*, **4119**, (San Diego), July 2000. Wavelet Applications in Signal and Image Processing VIII.
7. W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Trans. Patt. Anal. Mach. Intell.* **13**, pp. 891–906, Sept. 1991.

8. E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multi-scale transforms," *IEEE Trans. Inform. Theory* **38**, pp. 587–607, Mar. 1992.
9. N. G. Kingsbury, "Image processing with complex wavelets," *Phil. Trans. Royal Society London A* , Sept. 1999.
10. N. G. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Applied and Computational Harmonic Analysis* **10**, pp. 234–253, May 2002.
11. I. W. Selesnick, "Hilbert transform pairs of wavelet bases," *IEEE Signal Processing Letters* **8**, pp. 170–173, June 2001.
12. I. W. Selesnick, "The design of approximate Hilbert transform pairs of wavelet bases," *IEEE Trans. on Signal Processing* **50**, pp. 1144–1152, May 2002.
13. M. J. T. Smith and T. P. Barnwell III, "Exact reconstruction for tree-structured subband coders," *IEEE Trans. on Acoust., Speech, Signal Proc.* **34**, pp. 431–441, June 1986.
14. E. Dubois and S. Sabri, "Noise reduction in image sequences using motion compensated temporal filtering," *IEEE Transactions on Communications* **32**(7), pp. 826–831, 1984.
15. M. K. Ozkan, M. I. Sezan, and A. M. Tekalp, "Adaptive motion-compensated filtering of noisy image sequences," *IEEE Transactions on Circuits and Systems for Video Technology* **3**(4), pp. 277–290, 1993.
16. J. C. Brailean, R. P. Kleihorst, S. Efstratiadis, A. K. Katsaggelos, and R. L. Lagendijk, "Noise reduction filters for dynamic image sequences: A review," *Proc. IEEE* **83**(9), pp. 1272–1292, 1995.
17. B. C. Tom and A. K. Katsaggelos, "An iterative algorithm for improving the resolution of video sequences," in *Proc. SPIE: Visual communications and image processing*, **2727**, pp. 1430–1438, March 1996.
18. P. van Roosmalen, S. J. P. Westen, R. L. Lagendijk, and J. Biemond, "Noise reduction for image sequences using an oriented pyramid thresholding technique," in *Proceedings of ICIP-96*, **1**, pp. 375–378, IEEE Computer Society Press, (Lausanne, Switzerland), September 1996.
19. A. Kokaram, "Reconstruction of severely degraded image sequences," in *International Conference on Image Applications and Processing*, (Florence, Italy), 1997.
20. T. J. Burns, *A non-homogeneous wavelet multiresolution analysis and its application to the analysis of motion.* PhD thesis, Air Force Institute of Technology, 1993.
21. T. J. Burns, S. K. Rogers, D. W. Ruck, and M. E. Oxley, "Discrete, spatiotemporal, wavelet multiresolution analysis method for computing optical flow," *Optical Engineering* **33**, pp. 2236–2247, July 1994.
22. J.-L. Starck, E. J. Candés, and D. L. Donoho, "Very high quality image restoration by combining wavelets and curvelets," in *Wavelet Applications in Signal and Image Processing IX (Proc. SPIE 4478)*, 2001.